

PERFORMANCE MEASUREMENT PROBLEMS  
IN A PACKET-SWITCH NETWORK

Kevin L. Mills  
System Development Corporation

ABSTRACT

Central computing is slowly evolving to various forms of distributed computing based on computer communications network technology. These new computer network topologies will present significant problems to those charged with computer performance measurement responsibilities. This paper examines some problems faced when seeking performance measurements in a network based on packet switching technology. Specifically, the basic problems are: 1) what to measure and 2) how to measure. These problems are explored in the framework of a simple conceptual packet switch network.

1. INTRODUCTION

Over the past ten years, the concept of computer communications networks has moved out of the realm of research and development and into the public domain. The key technique that was brought to public usage was the packet-switch network. Most of the research and development has been concentrated on the communications vehicle provided by the packet-switching technique. For this reason, a number of problems remain extant when examining the performance of the entire computer communications network rather than just the communications portion of the network.

The purpose of this paper is to illustrate some problems in measuring performance in a packet-switch network. Specifically, two questions are examined:

- 1) What to measure?, and
- 2) How to measure?

These questions are discussed in three major sections as follows.

The first section, "A Simple, Conceptual Packet-Switch network," describes a simple, non-existent packet-switch network at a conceptual level. This description provides the reader with a frame of reference for the following discussions.

Section two, "What to Measure?", investigates selected problems inherent in defining performance measures for the network presented in the preceding section. Existing literature references are cited to present the views of individuals experienced in packet-switch networks. The major aims of this section are to identify problems which exist in defining performance measures and to provide insight into why the problems exist.

The third section, "How to Measure?", investigates problems related with collecting data required to measure performance in a packet-switch network. Of course, data collection problems are largely a function of the measurement definition problems identified in the previous section; however, measurement techniques would remain a topic worthy of separate treatment even if all performance measures were well-defined and universally accepted.

The paper ends with "Conclusions" providing a summary of the main points presented in the preceding sections. Also, some research directions are offered to help solve the problems discussed in this paper.

2. A SIMPLE, CONCEPTUAL PACKET-SWITCH NETWORK

There are a number of packet-switch networks operating today. These include: TELENET [HOV76], TYMNET [SMI78], Packetnet [WEC76], ARPANET [MAR76, KAH77], and WIN [WIN74]. There are also several planned for the near future.

These include: A T & T's ACS [DAV79] and the Department of Defense's AUTODIN II [AUT77].

Each existing or planned packet-switch network differs in the implementation of the concept. For this reason, the discussion presented in this paper will be oriented toward a nonexistent general packet-switch network that presents the major functions in such a network. Because the first major packet-switch network was ARPANET, most of the concepts presented in the network described below were derived from ARPANET and its derivatives WIN and TELENET.

The conceptual packet-switch network will be divided into three major components: 1) the communications subnet, 2) the hosts, and 3) the terminals. Each of these major components is described below.

2.1 COMMUNICATIONS SUBNET

The communications subnet in a packet-switch network provides the communications vehicle for the geographically distributed hosts and terminals. Figure 1 illustrates the communications subnet for the packet-switch network being discussed presently. The communications subnet is made up of communications facilities, packet-switch nodes, and a network monitor center.

2.1.1 COMMUNICATIONS FACILITIES. Every packet-switch node is connected to two others by 56 Kbps leased lines. Not shown in the diagram are modems (or data service units) which are present at each end of every leased line (the lines may be analog or digital). The purpose of the leased lines and modems is to provide full-duplex communications between two packet-switch nodes.

2.1.2 PACKET-SWITCH NODES. There are four packet-switch nodes. Each node is a minicomputer. Three types of nodes exist. A host packet-switch node can communicate with up to four host computers and two leased lines. These nodes are numbers 1 and 2 in Figure 1. (NOTE: this is a simple network compared with ARPANET which contains over 100 packet-switch nodes.)

A second type of packet-switch node is a terminal node. Terminal nodes can communicate with up to 32 terminals and two leased lines. Node number 3 in Figure 1 represents a terminal node.

Both the terminal and host packet-switch computers have a TTY console and no secondary storage as illustrated in Figure 1. These packet-switch computers provide physical link control facilities for the high speed leased lines and the entry connections to the hosts and/or terminals.

The third type of packet-switch node is the network monitor packet-switch. The network monitor node provides the same functionality as the other nodes

Figure 1.

# A SMALL (FOUR NODE) PACKET SWITCHED COMMUNICATIONS SUBNET

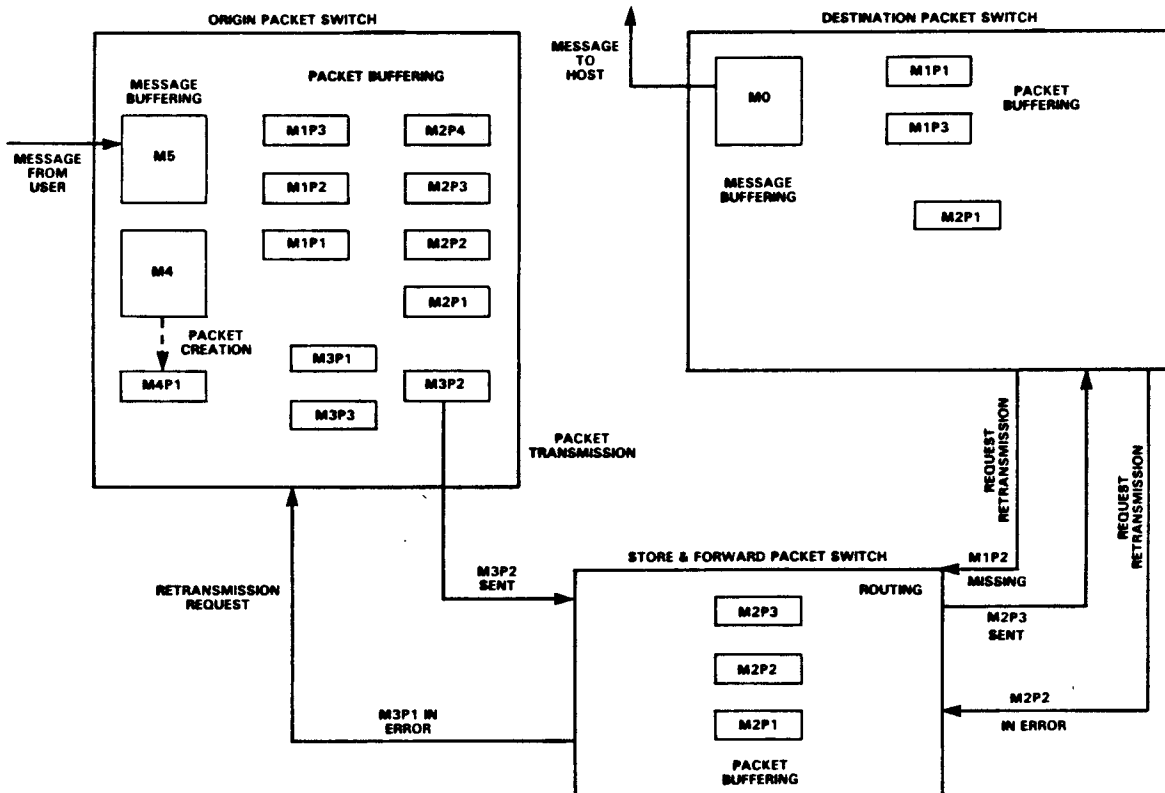
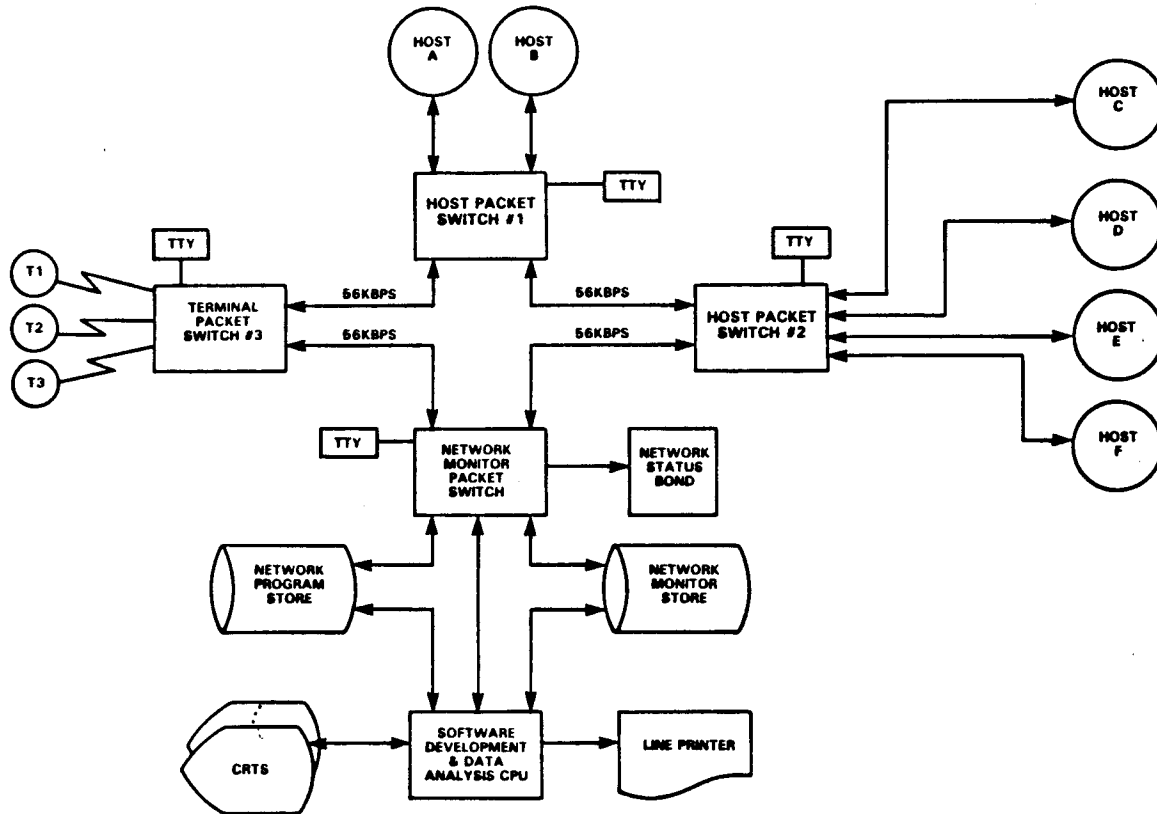


Figure 2.

## ROLES OF A PACKET SWITCH COMPUTER

with the following important additions. The network monitor packet-switch drives a network status board display and has access to two types of secondary storage. Storage is provided for network programs which can be loaded into other network packet-switch nodes from the network monitor center. Storage is also provided for live recording of network performance data. The network monitor packet-switch is combined with other components to form the network monitor center.

**2.1.3 NETWORK MONITOR CENTER.** In addition to the packet-switch, the network monitor center contains a large scale mincomputer to do network software development and data analysis. This computer shares access to the network program storage and performance data storage with the packet-switch, allowing off-line software development and data reduction and analysis.

The network status board is also a part of the network monitor center. The status board provides a visual presentation of the status of all leased lines, modems, packet-switch computers, and host computers throughout the network.

Physically, the network monitor computer and the network monitor packet-switch are directly connected. This allows the network monitor to become a host on the network, if desired.

**2.1.4 COMMUNICATIONS SUBNET OPERATION.** Operation of the communications subnet is the primary function of the packet-switch computers. A packet-switch computer receives a message from a terminal or host connected to it. The packet-switch must identify the destination of the message, divide the message into one or more packets, and send each packet over the appropriate leased line toward its destination. In a multi-packet message, each packet can take a separate route, and therefore, can arrive at the destination packet-switch out of order.

A second major job of a packet-switch computer is to assemble all packets destined for a host or terminal connected to it into a message that is properly ordered and to send that message on to its host/terminal when it is completely assembled. Associated with this responsibility is the task of detecting missing packets and requesting the retransmission of these packets by the originating packet-switch node. Duplicate packets must also be detected and discarded.

The third primary job of a packet-switch is to act as a store-and-forward node for packets that are passing through enroute to another packet-switch computer. The only responsibility that a forwarding packet-switch has for these packets is to check for transmission errors and to request retransmission if the packet was in error.

These three major roles of a packet-switch computer are illustrated in Figure 2. Packet-switch computers also perform many other functions that are beyond the scope of this discussion.

## 2.2 HOSTS

Hosts in this network are represented by large computer systems such as: IBM 370, UNIVAC 1108, Honeywell 6000, or DECSYSTEM 10. The network has six host computer systems (A - F in Figure 1) and the hosts are heterogeneous (i.e., there are several different vendors' computers connected to the network).

Software executes in each host to allow access to the network for the purpose of establishing logical connections to other network hosts. This software will be called the Network Access Program (NAP). The NAP serves as the basis for accomplishing several types of

network-oriented tasks. These tasks are: 1) local terminal access to remote hosts, 2) teleconferencing, 3) file transfer, and 4) workload sharing. Each of these tasks will be discussed below. Note that each host has some configuration of online secondary storage, local terminals, and off-line auxiliary storage. The amounts of each type of resource may differ for each host.

**2.2.1 LOCAL TERMINAL ACCESS TO REMOTE HOSTS.** Each host supports a number of local terminal users. Software executing in each local host allows local terminals to establish connections and to operate as if each user terminal is local to a user selected remote host. This also requires that each host be able to accept messages from remote terminals as if they were local.

A major problem in operating as described above is that the several hosts and terminals involved may operate at different speeds, with different line protocols, and with different character code sets. To alleviate this problem, each host must provide software to convert its unique local terminal characteristics into a network virtual terminal and to convert network virtual terminal characteristics into its unique local terminal requirements. The network virtual terminal is just an agreed upon set of terminal characteristics so that any host interfacing to the network can participate in the remote terminal access. In the network described here, the network virtual terminal is a 300 bits per second, asynchronous, ASCII code terminal.

To gain access to a remote host, a terminal user must have a valid account at the remote host. Once a terminal gains access to the remote host, the terminal user must operate with the command language appropriate to the remote host.

**2.2.2 TELECONFERENCING.** Teleconferencing is an extension of the terminal access to remote hosts. In teleconferencing, a conference file is established at a host and a number of terminals are allowed simultaneous access to the contents of the file. A high level protocol is established to ensure that only one terminal user can write to the file at any instant, while all other users can view the file simultaneously.

This function provides a means of written conferencing among geographically separate users who have access to a terminal on the network. The high level protocol is established so that the teleconferencing user does not have to be aware of what computer is maintaining the conference file. However, a complex and time consuming procedure is necessary to establish the teleconference resources prior to the beginning of the teleconferencing session.

**2.2.3 FILE TRANSFER.** This capability allows terminal or batch users on any network host to transmit files between any set of hosts, assuming valid accounts on all hosts involved. A command language is established to facilitate this transfer, but some of the statements are host dependent.

When data is transferred through the network, it is in a common network file transfer format. For this reason, each host must support the conversion to and from the network file transfer format. Also, each host must be able to support a checkpointing system when it is involved in a file transfer. This checkpointing system is designed to avoid the great expense of total retransmissions when a network component fails during a large file transfer.

**2.2.4 WORKLOAD SHARING.** Workload sharing is a special extension of file transfer. In the network under consideration, workload sharing is only meaningful between two network hosts of the same vendor. A workload sharing operation consists of transferring needed command

language, program, and data files to a remote host for execution. The resulting output files are then transferred back to the originating host for disposition.

This service may be useful when the two hosts involved are in widely different time zones. Under this circumstance, the non-peak hours of one host can be used to provide supporting service during the peak workload hours of other hosts.

To use this service, special accounting procedures have to be set up so that the use of resources can be correctly assigned. In the network under consideration, a special accounting information file is sent to the remote host at the start of a workload sharing period and a resource utilization file is returned to the originating host at the end of the period.

### 2.3 TERMINALS

Terminals access the network from two general locations. As already described, terminals local to a network host can use the network to access remote hosts. The second way a terminal can access the network is through the terminal packet-switch node (#3 in Figure 1).

Terminals accessing the network from the terminal packet-switch node can be of any type that uses an asynchronous line protocol. Each terminal has the necessary characteristics converted to the network virtual terminal by the packet-switch computer. The packet-switch computer also converts messages intended for each terminal from network form back to the unique requirements of that terminal.

Terminals accessing the network can be video or hard copy devices with or without internal memory. Generally, the network strives to provide support for most types of terminals.

### WHAT TO MEASURE

This section investigates some of the problems inherent in defining performance measures for a packet-switch network. The discussion is organized around the three major components of the network: 1) communications subnet, 2) hosts, and 3) terminals. Performance measures will be presented that are relevant for each of these components. Finally, the shortcomings of the performance measures will be discussed.

#### 3.1 COMMUNICATIONS SUBNET

Four performance measures are cited overwhelmingly for communications subnets: 1) message response-time [TOB78, KLE78, NAC77, BRA78A, BRA78B], 2) throughput [TOB78, KLE78, NAC77], 3) reliability [KLE78, NAC77, TOB78], and 4) cost [KLE78, NAC77]. Unfortunately, these performance measures are not always defined in the same manner.

In [NAC77], message response-time is described as a function of the application and where the response is to be measured. Two types of response-times are cited: 1) terminal response-time (from the time the user presses the send key until he receives the first character of output), and 2) overall response-time (from the time the user generates a message until he receives a complete reply).

In [TOB78], message response-time is described as the round-trip delay from the time a message is sent until a request for the next message is received. This definition is a subset of overall response-time defined in the preceding paragraph.

In [BRA78A and BRA78B], message delivery-time is seen as more important than actual response-time or throughput. The point to be made here is that the user is primarily

concerned with the speed with which his messages arrive at their destination. Since the network is constructed so that no one physical path is dedicated for sending messages between a source and destination, one-way message delivery-time is much more meaningful than message response-time.

Construction of the communications subnet so that dedicated physical paths do not exist is also addressed in [KLE78]. The author points out that when using a packet network a user cannot get the same quantitative values for performance measures as he can with a dedicated, leased line. Moreover, the quantities can seldom be measured in a straightforward manner (this is further discussed in Section 4).

The second commonly cited performance measure in the communications subnet is throughput. [TOB78] looks at throughput from the viewpoint of destinations. In this view, throughput is the sum of all packets received at their destination successfully (per unit time). This excludes all duplicate packets generated because of missed acknowledgements or transmission errors.

In [NAC77], throughput is described in terms of capacity as the maximum traffic a system can carry with a specified response-time. It is not made clear what measures will define this capacity.

[KLE78] states that throughput and capacity of a packet network are not straightforward concepts. Although throughput can be defined as messages correctly received at the destination, this would depend on the traffic pattern that is assumed.

It is clear that, although not everyone agrees on the definition of throughput for the communications subnet, messages that are not correctly received are not valid contributors to throughput measures. This poses a difficult problem because packets that must be retransmitted have used some amount of system resources yet do not count as network throughput. But these erroneously transmitted packets will still have an effect on message delivery-time. This problem is addressed in some depth in [BRA78A, BRA78B].

A third measure of communications subnet performance often cited is reliability. Generally, two ways of looking at reliability are extant. First, reliability can be defined as the percentage of time a user can communicate with required points in the network. This is the system-wide view previously described in [NAC77].

An alternative view is to examine reliability on a component basis. Each component is defined as a function of its probability of failure and its impact on the network if it does fail. This is an attempt to derive quantitative measures of network reliability as probabilities. This approach is discussed in [TOB78].

The fourth generally cited communications subnet measure is cost. Though often cited, the relationship of cost to other performance measures is rarely defined. The difficulty with cost is that its significance depends on who you are with relation to the network. If you are using someone else's network, then you would evaluate their charges versus the cost of alternatives providing the same levels of performance. If you are the owner of a private network, you might examine the costs of remaining closed versus renting the use of your network to outsiders for a fee. Essentially, cost is a complicated economic issue that the technical writers tend to cite only in passing.

Although the above performance measures are most often discussed, many others are found in the literature. It is beyond the scope of this paper to discuss these less

often cited measures of performance, but they will be mentioned and references are provided.

[NAC77] cites network sensitivity to peak load traffic, packet route blocking probability and transmission error rates as other communications network performance measures. [TOB78] indicates that the basic performance measures have some limitations and must be augmented by special measures such as: counts of the number of times a packet is caught in a network loop (this is a complement to throughput), measures of fairness in allocating network capacity among large and small users, and measures of network congestion.

Other performance measures cited include: availability, generality, ease of use [GIL78], expandibility [BEC78], maintainability [LES78], and probability of message misdelivery [AUT77]. One general observation is that those measures of communication subnet performance that are most easily quantifiable and most specific have not yet been successfully related in a general model of network performance.

### 3.2 HOSTS

Hosts are the components of packet-switch networks that have been given the least attention in assessing network performance. Much research has been conducted relative to evaluating the performance of a mainframe computer in a stand-alone mode, but the added problems of host performance evaluation in a network environment have not been addressed to any great extent.

Some questions that might be important are: who is using how much of what resources how often and for how long? and, what impacts are network operations having on a host's local operations?

A network member host must run a network access program. This program is sheer overhead, providing only a path to other network nodes. To accomplish useful network activity, other programs must also execute in the host. These can include user and server remote terminal access programs, file transfer programs, remote job entry programs, workload sharing programs, and teleconferencing programs [BEN74].

Performance goals might be established for each of these services. For example, it may be desirable to keep terminal response-time the same for remote and local users or to keep the load placed on the host by the file transfer program low enough so as not to affect local data base operations [BEN74].

Since remote terminals use more resources than local terminals, statistics to indicate how the increasing support for remote terminals affects support for local processing might be desirable [GIL78]. Performance measures showing how workload sharing operations affect job throughput and turnaround-time might also be useful [BEC78].

Other measures of performance may be required to show how the host is affected by interactions with the communications subnet. For example, as delays increase in the communications subnet, the packet-switches may begin to reject further input of messages from the hosts. This will cause the messages to be queued by the host outside of the communications subnet [KAH77]. Host measures might be needed to assess the impact of this message queuing on the performance of the host.

Further measures of performance might be required to show how one network host is affected by interactions with other hosts on the network. For example, a given receiving host may be unable to accept all messages at their arrival rate, or may fail in the middle of a file transfer, or may simply discard valid messages [KAH77].

Additional performance measures would aid in describing how these occurrences affect the transmitting hosts.

Generally, performance measures for the operation of a host computer in a network environment are not well defined. The problems to be overcome before such measures can be established are complex. It is clear that these measures cannot be established by viewing the host in isolation from the other network components.

### 3.3 TERMINALS

When assessing network performance at the remote terminal positions, the primary emphasis is on the end user. Not much work has been conducted to assess network performance from the terminal user's perspective; however, the National Bureau of Standards [ABR76, ABR75] has been investigating this approach and Philip Elam of CINCOM [ELA78, ELA79] is a strong advocate of this approach.

Researchers at the National Bureau of Standards (NBS) feel that performance "...measurement techniques that indicate percent utilization of a CPU or some other internal measure of performance are comparatively meaningless to a remote user" [ABR76, p3]. Although no measures of network performance have been proven meaningful to users, the NBS researchers offer some suggestions for further examination.

The network performance measures suggested by NBS can be placed in four classes:

- 1) Time-based Measures,
- 2) Measures of Length,
- 3) Measures of Multiplicity, and
- 4) Rates.

Each of these will be discussed below.

Time-based measures serve to provide a profile of how the terminal user's time is spent during a terminal session (from logon to logoff). Time measures to be made include: system response-time as a measure of request to fulfillment, system print out time as a measure of verbosity, acknowledgement delay as a measure of system awareness of user requests, length of time until acknowledgement to the terminal that the system is aware of a user request, user think time, user typing time, the time it takes a user's task to be completed by the system, and the time it takes the user to complete a total session.

The significance of these measures is that they not only measure the system's response to the user but the user's response to the system. Measuring the user's response to the system can illustrate such problems as an overly verbose system/user dialogue or an unnecessarily complex command language.

The second class of measures, length, serve to provide a profile of how much of the system network resources are being used during a terminal session. Measures of length include: the number of characters transferred between system and user broken into user messages, system responses, and system overhead. Another measure of length is the number of transactions that comprise user tasks and the entire session. These measures are primarily aimed at isolating the communication costs associated with user terminal sessions.

The third class of user-oriented network performance measures cited by the NBS researchers is measures of multiplicity. These measures view the user's terminal session on a logical level. Measures include: the number of times each command is used, the number of times each task is used, the number of different processors used (these could be programs or actual CPUs), the number of times each processor is called,

the number of times each system and user error occurs, and the number and identity of faulty transactions. The information to be derived from these measures could also serve as input to evaluate the performance of network hosts and to plan network expansion based upon actual use of the network.

fourth class of measures to be considered is rates. These measures are generally derived by evaluating some of the above quantities with respect to time. Rates include such items as: character arrival and departure per second at the terminal, transactions per hour, tasks per hour, and errors per hour. Other ratios are derived by looking at a class of items as a percentage of a total. Some examples include: correct transactions divided by total transactions per session as a measure of reliability; number of each type of command used as a percent of total commands used; percent of total time that the user is waiting; and an interference ratio calculated as the ratio of actual response-time to optimal response-time. These measures provide logical utilization and throughput figures that may be more akin to a user's understanding than are hardware utilization and throughput measures.

The scope of measures contained in the above four categories is very broad. Because a network may provide service to several hundreds of terminal users, a major question arises, "how can all these measurements be made for each terminal user?" This question is addressed in Section 4 of this paper.

In contrast to the approach taken at NBS, Philip Elam adopts a nonquantitative approach to network performance evaluation from the terminal user's perspective. Mr. Elam believes that a network's ultimate performance is determined by how effectively it is operated by the user community. Under this view, users must make use of the network in order for it to have any economic benefits. Therefore, Mr. Elam places primary importance on those characteristics of a network that determine how often and how effectively the network is used.

For Mr. Elam, the overriding network performance measure is ease of use. This measure is derived primarily from Mooers Law [ELA78, p. 51]:

An information retrieval system will tend not to be used whenever it is more painful and troublesome for a user to have information than for him not to have it.

Unfortunately, ease of use is difficult to quantify.

Mr. Elam identifies some items that might be considered when determining ease of use. These include: response-time predictability, keyboard layout, command language, network reliability, and effective error messages. These items still present difficulty in obtaining quantitative values.

Mr. Elam suggests that the key is to know the user. Know what the user does, what he likes, what he dislikes, and how he thinks. Does the user feel pressured by rapid response-time? Does the user fear that mistakes made at his terminal will cause system failure? Does the user feel that the system has no sympathy for him? Does the user hate to type?

Only after a user profile is accumulated can a quality decision be made as to what performance goals are desirable. Mr. Elam suggests that the evaluation of a network be based on how well it returns relevant data to the user. This includes such factors as: timeliness, precision, readability, and ease of access. But rather than assigning strict performance goals to these factors, the user must be understood and goals established and varied to satisfy his needs.

### 3.4 PROBLEMS

The problems that arise when defining performance measures in a packet-switch network are caused primarily by the scope and complexity of the system being considered. Many problems exist when measuring performance in a single, large-scale computer system. When several of these systems are connected via a new network technology, problems that already exist increase exponentially and new problems arise.

In a communications subnet, there are thousands of hardware and software variables. Each of these variables is tunable and their interactions can have impact on throughput and message response-time. Such variables as disk accesses per second, core space available, CPU speed, CPU instruction set, communications equipment, I/O channel utilization, program design, and queuing methodology cannot begin to be related in a model that has predictively valid results [AUD78]. Even the performance measures these values affect have their own independent relationships (e.g., as throughput increases, response-time decreases). With these general thoughts as background, some specific problems will be addressed.

From the discussion in Section 3.1 above, it is clear that no performance measures are precisely defined and widely agreed upon for a communications subnet. Since messages go out as packets and packets can take various physical routes to the same destination, response-time can be viewed from various angles. None of these has been shown clearly preferable to any other.

Throughput in the communications subnet is clouded by the issue of packets that become lost along the way and must be retransmitted. Everyone agrees that these do not count as valid packets in network throughput evaluation, but they have a deleterious effect on message delivery-time.

Other performance measures in the communications subnet are not quantifiable in a meaningful way. That is, a clear relationship between those measures and user performance and satisfaction cannot be demonstrated. This results primarily from the complexity of the whole network.

For host computers on the network, the same performance measures that have always been used might be meaningful with the provision for remote vs. local statistics. The actual use of hosts in a packet-switch network has not been on-going long enough for any real progress to be made in defining meaningful measures of performance.

The most promising approach to evaluating packet-switch network performance seems to be based on the terminal user. Philip Elam's contention, that network performance evaluation should be based on the fact that a network exists to serve the end user, provides a solid starting place for examining network performance. The NBS has suggested numerous measures that may help in this examination. What has not been shown is which of those measures reflect facets of network operation that are important to the user. The NBS has suggested relationships between their measures and user criteria, but more work must be done to prove these relationships.

From the preceding discussion, it can be seen that problems exist when evaluating network performance relative to each of the major network components. But a much more difficult problem, performance of the network as a whole, has not been addressed. When this broad view is taken, relationships may be found that show that optimal performance for one, two, or all the major components may have a deleterious effect on the network as a unit. Much more effort is required to

develop performance variables and establish relationships that provide a meaningful picture of overall network performance as seen by the network user.

#### 4. HOW TO MEASURE?

This section discusses the problem of how to make performance measurements in a packet-switch network. The issue is addressed independently of the previous issue; that is, an assumption is made that the performance measures are defined. (The author realizes that the two issues are probably not independent. However, to facilitate the organization of this paper, they are treated here separately.) Since each of the major network components may require different measurement techniques, the discussion is organized into three sections examining the communications subnet, hosts, and terminals. The problems associated with the measurement techniques will then be summarized.

##### 4.1 COMMUNICATIONS SUBNET

The communications subnet (as presented here) has three categories of built-in software measurement tools:

- 1) Packet-Switch Background Programs,
- 2) Network Monitor Center Programs, and
- 3) Data Analysis Programs.

These software tools are also found on the ARPANET communications subnet. Since the example conceptual packet-switch network is based upon the ARPANET, the description of these tools is taken from [POH75]. The conceptual packet-switch network might also contain hardware tools as discussed in 4.1.4 below.

**4.1.1 PACKET-SWITCH BACKGROUND PROGRAMS.** These programs run in each packet-switch as background programs. They are implemented as pseudo network hosts. Therefore each background program has a unique network address and certain types of network messages are sent and received by these pseudo hosts as statistics messages.

One function performed by the background programs is a trace that allows the progress of messages to be tracked through the network packet-switches. This trace is useful for studying the routing, queuing, and packet-switch processing aspects of the network. (NOTE: these programs were built primarily as a system development tool.) The information traced includes: time the last bit of packet arrives at a switch, time the packet was put on an ongoing queue, time the switch started to transmit the packet, time an acknowledgement was received for the packet, the length of the packet, the destination address, and the packet-switch header. All of the accumulated data is sent as a trace message across the network to the Network Monitor Center for printing and/or storage.

A second function performed by the background programs is the generation and accumulation of statistical data. Each packet-switch is capable of generating a snapshot of its condition and sending that information in a message to the Network Monitor Center. This information includes length of queues in the switch, the current routing table for the switch, and the present delay table.

Besides snapshots, each packet-switch can collect cumulative statistics reflecting its operation. These statistics are accumulated in the packet-switch memory and sent to the Network Monitor Center at predirected intervals. These statistics describe message characteristics, time intervals, and channel usage. Message characteristics include: message lengths entering and exiting the network at the packet-switch, length of the last message that entered and left the network, total number of words in all lost packets, total number

of messages from all hosts attached to the packet-switch, and total number of control messages to each host. Time statistics include: the sum of all round trip times to each destination, total number of round trips to each destination, and round trip time overflows. Channel statistics include: number of link maintenance messages sent per channel, number of data words sent and received per channel, number of errors detected per channel, and number of times buffer areas were unused or overflowed per channel. Most of these statistics are measurements of message delivery-time, packet-switch throughput, and network overhead.

A third function supported by the background programs is test message generation. A program builds and sends artificial, fixed-length messages to a destination at a fixed interdeparture time. A program at the destination discards these artificial messages. This artificial load can be generated by the communications subnet to check its own operation in the absence of legitimate user messages.

The final function performed by the background programs is the generation of two kinds of status reports. First, a packet-switch to Network Monitor Center status report is sent every 52 seconds or upon change of status. This report includes the status of hosts, lines, and memory attached to the packet-switch. The second report is a throughput report sent to the Network Monitor Center every 52 seconds. This report contains the number of packets and words transmitted during the last minute by each line and the number of messages, packets, and words in and out of the packet-switch from each host.

The operation of all of the above background programs is controllable at each packet-switch. An operator can inspect and change the relevant parameters using the TTY at each node. The Network Monitor Center can also change the relevant parameters by sending command messages to the appropriate packet-switches.

**4.1.2 NETWORK MONITOR CENTER PROGRAMS.** The Network Monitor Center (NMC) programs execute in the network monitor packet-switch (see Figure 1) and serve two primary purposes: 1) to collect, accumulate, and record statistics created by the packet-switch background programs, and 2) to monitor continuously and display the status of all network components. In support of the first purpose, the NMC programs record throughput reports sent from the network packet-switches. The NMC maintains throughput tables for each host and line in the network and updates these tables at preset intervals. Each hour, the accumulated line and host throughput tables are dumped to secondary storage. Every eight hours, these statistics are reported. Finally, every twenty-four hours a summary of all packets transferred through the network is printed and the tables are initialized to begin a new day.

To support the second purpose, the NMC programs are augmented by a network status board. The status board presents a visual representation of the hosts, lines, and packet-switches in the network. Each packet-switch is programmed to examine itself and its environment periodically and to report the results to the NMC. The NMC programs accumulate a global picture and store this information while providing a summary for the operator. Significant changes in status are reflected on the network status board accompanied by an audible alarm. This allows the operator to quickly identify failures detected by the NMC programs.

**4.1.3 DATA ANALYSIS PROGRAMS.** These programs execute in an off-line computer (see Figure 1) using the data written to secondary storage by the NMC programs. The data analysis programs allow statistics to be analyzed and reported in meaningful ways without interfering

with network operations. Additionally, the data analysis programs can be set up as pseudo network hosts so that statistics can be reported directly without being accumulated by NMC programs.

illustrate the increased capability provided by the analysis programs, a list of statistics that can be calculated follows: message and packet size distribution, mean round trip delay, mean traffic-weighted path length, flow of traffic to and from hosts connected to the same packet-switch, identities of the most heavily used packet-switches and lines, hosts that send most of their messages to a limited number of destinations, and utilization for each line.

**4.1.4 HARDWARE TOOLS.** Many hardware tools are available that attach to communications lines at various points and monitor traffic flow through these points. Some of these will be examined in Section 4.3. The discussion here concerns the incorporation of microprocessors as an integral part of the communications subnet.

With the decreasing cost of microprocessors, it is becoming increasingly feasible to include microprocessors in every modem and other communication component and to provide a central minicomputer to poll each microprocessor on a scheduled basis. This arrangement has been proposed as the backbone for a network management system [ARM79]. The significant advantage is that this kind of monitoring system can operate independently of the communications subnet that it is monitoring and, therefore, reduce measurement artifact. Since microcomputers are programmable, many of the same statistics discussed in the preceding three sections could be accumulated. Further, the microcomputer could be programmed to take corrective actions when failures are noticed.

## 2 HOSTS

Presently, no special mechanisms are available for accumulating network related information at host sites. S. Poh suggests that the network access program in each host is a likely place to collect such information [POH75]. Several reasons exist for this lack of progress. First, since packet-switching technology is fairly new and expensive to implement, all development has tended to involve shared use of the packet-switch network among users who have very different goals. Therefore, development of host specific performance measurement tools has been left to each host site. These sites typically have had trouble implementing the network access program and network application programs; consequently, they have seldom addressed the problem of network performance.

Secondly, the network technology has typically been imposed upon an already existing host operating system that was not designed for this kind of network application. Therefore, no efficient ways exist to incorporate appropriate network performance monitoring software into the operating system.

What is usually seen when measuring the performance at network hosts today is the allocation of special job numbers to network jobs. Then the regular job accounting packages can be used to report on network jobs in the usual manner. The analyst at the host site can then identify the network jobs and calculate statistics for these jobs separately from the local jobs. A more desirable approach is to incorporate the needed tools into an operating system that is designed to support computer operations in a network environment.

## 4.3 TERMINALS

The measurement of network performance at terminal locations is a difficult problem because of the large

number of terminals typically found in a network and the lack of processing resources usually available at each terminal. One commonly used approach is line monitoring. In this approach, a line monitor is attached between the terminal and modem (or other communication device). The line monitor can then display and record data and control characters that pass by on the line [STU78]. Some line monitors also allow the generation of test messages to check out the terminal's response to polling requests [LAP76]. Most of these devices are primarily used as diagnostic tools rather than performance measuring tools. They are general and simple so that they can appeal to the widest possible market.

An extended form of this approach is marketed by Tesdata Systems Corporation [TES80]. The Tesdata MS-109 replaces the line monitor with a line interface module (LIM). The LIM is programmed to recognize the protocol that is being used on the line and is also programmed to collect meaningful performance data (e.g., response time, utilization, and availability). Data can be accessed via a video terminal or can be data-based on a larger machine for later reporting. This approach works well for small to medium star networks (up to 256 lines) but, in its present form, is too limited for use collecting the kind of measures suggested by NBS researchers.

Recalling the performance measures suggested by NBS researchers (see Section 3.3), it is evident that line monitors can collect the kind of data required but they cannot manipulate that data into useful information. To overcome this problem, the NBS used a Network Measurement Machine (NMM) in its research [ABR76, ABR75]. The NMM was attached to a communications line in the same way as a line monitor. However, the NMM was composed of a minicomputer system that provides significant processing power. With this technique, meaningful information can be generated at useful rates and reported in a summary fashion at the monitoring point. This is similar to the Tesdata approach except the type and numbers of measures is greatly expanded.

The problem with the NMM is cost. To provide a minicomputer system at each terminal site would require substantial cost/benefit justification. However, the potential for this approach becomes attractive as the availability of cheap processing increases. With more intelligence available at every terminal, the NMM functions could be made a part of the terminal's logic. Or, using microcomputers, dedicated NMM hardware may become cost-effective. Generally, this approach seems to be increasingly feasible as shown by the emergence of commercial products such as the Tesdata MS-109.

## 4.4 PROBLEMS

Because a packet-switch network is a complex and expensive system, much of the burden has been shared by the groups who jointly developed the existing networks. With differing goals and research directions, an inconsistent and unsatisfactory approach to the development of network performance measurement techniques has resulted.

Most of the work has been concentrated on the communications subnet. This is understandable because the subnet represents the new technology. The major problem with the current techniques for measuring performance in the communications subnet is the large overhead they incur. The subnet becomes flooded with messages reporting status and throughput statistics to the NMC programs, and as user traffic increases the network capacity taken up by these messages becomes more costly.

Probably the most pervasive problem is the lack of provision for coordinating and analyzing the performance measures taken at terminals, in the subnet, and at the hosts. When all of these statistics can be accumulated and related at some central point, a true picture of total network performance will be possible. To achieve this objective, a dedicated measurement network may have to be constructed for operation in parallel with the network being measured. This is an expensive solution, but one which may become increasingly feasible as the cost of processing logic decreases. The concept appears worthy of investigation.

If found to be a useful approach, the problem then will be to decide who is responsible for implementing such a solution. In an environment where the majority of networks will be built for subscriber service, the cost of performance measuring will be passed off to the subscribers. The subscribers must bear the responsibility for making their performance measurement requirements known in the near future to companies building such networks. Otherwise, networks will continue to be constructed with the primary performance measurement tools aimed at the communications subnet and, thereby, leave the network user with no tools for making performance measurements meaningful to him.

## 5. CONCLUSIONS

This paper has examined performance measurements in a packet-switch network composed of a communications subnet, hosts, and terminals. The bulk of research done in performance measurement for this type of network has been aimed strictly at the communications subnet. Very little research has been conducted to evaluate the role of hosts in network performance. Some research has been conducted at NBS to examine network performance from the terminal user's viewpoint.

Within each component, performance measures are not well defined or widely agreed upon. Further, no research has been done to relate performance measures for the communications subnet, hosts, and terminals together into a total performance measurement model for a packet-switch network. Attempts to create this model would prove valuable.

In the realm of data collection techniques, some well defined means are available for each component of the network, but problems exist with each technique. Communications subnet measurements are made by transferring status messages through the net, thus affecting the very performance they are intended to measure. Host data collection relies on operating system accounting packages and procedural separation of data. Terminal measurements have been taken by tapping the line between the terminal and the communications subnet with minicomputers and microcomputers. This approach is, presently, too expensive to be used at every terminal location on a large network.

The principal problem with present available performance measuring techniques is that they are isolated from one another. To support measurements made in accordance with the total network model that has yet to be developed, a means of coordinating the measurements taken from all three network components will be required.

The approach that shows perhaps the most promise is to examine the network performance problems from a terminal user's perspective. With the decreasing cost of microprocessors and the increasing communications bandwidth available from digital communications and fiber optics, a performance measurement network operating in parallel with the operational network may well become feasible. Such a measurement effort can attack the total problem without interfering with the operational network. This approach can also focus attention

on the primary performance evaluator, the network user. A philosophical shift of this nature will support determination of what constitutes optimum performance for the terminal users on the network and allow adjustment of the network operating variables so that they support the user's view of optimum performance.

## 6. ACKNOWLEDGEMENTS

The author wishes to thank Charles Ross of Tesdata Systems Corporation for his encouragement, Paul Gruber and Mike Williams of System Development Corporation for their timely and insightful comments, and Dr. Edward Ries for his financial support. A special thanks also goes to William Bowman of System Development Corporation for introducing the author to a real packet-switch network.

## REFERENCES

- ABR76 Abrams, Marshall D.; Treu, Siegfried; and Blanc, Robert P. Measurement of Computer Communication Networks. Washington, D.C. : U.S. Department of Commerce (NBS), July, 1976.
- ABR75 Abrams, M. D. and Cotton, I. The Service Concept Applied to Computer Networks. Washington, D.C. : U.S. Department of Commerce (NBS), August 1975.
- ARM79 Armstrong, Dr. Thomas R. "An Automated Network Management System." Mini-Micro Systems, March, 1979, p. 79.
- AUD78 Audin, Gary. "Performance: The Battle of Needs Vs. Desires." Data Communications, October, 1978, p. 99.
- AUT77 System Performance for AUTODIN II. Washington, D.C. : Defense Communications Engineering Center, February, 1977.
- BEC78 Becker, Hal B. "Distributed Data Processing: Where We Are Now." Data Communications, August, 1978, p. 77.
- BEN74 Benoit, J. W. and Wood, D. C. Network Front End Study. McLean, VA : Mitre Corp., August, 1974.
- BRA78A Brayer, Kenneth. "Time, Not Throughput, For Net Designs: Part I: Techniques." Data Communications, October, 1978, p. 39.
- BRA78B Brayer, Kenneth. "Time, Not Throughput, For Net Designs: Part II: Applications." Data Communications, November, 1978, p. 83.
- BUT77 Butto, M., et al. Models for the Performance Evaluation of a Packet Switching Network with Retransmission Time-out. National Telecommunications Conference, 1977.
- DAV79 Davis, George R. "AT&T Answers 15 Questions about its Planned Service." Data Communications, February, 1979, p. 41.
- ELA79 Elam, Philip G. "Network Efficiency Takes More than Speed Alone." Data Communications, March 1979, p. 103.
- ELA78 Elam, Philip G. "Considering Human Needs can Boost Network Efficiency." Data Communications, October, 1978, p. 51.
- GIL78 Giles, Howard L. "Successful Network Management Hinges on Control." Data Communications, August, 1978, p. 33.

- HOV76 Hovey, Richard B. "The User's Role in Connecting to a Value-Added Network." Basics of Data Communications. Edited by Harry Karp. New York: McGraw-Hill, Inc., 1976.
- KAH77 Kahn, Robert E. "Resource-Sharing Computer Communications Networks." Tutorial: Distributed Processing. Edited by Dr. John Berkowitz. New York: IEEE, 1977.
- KLE78 Kleinrock, Leonard. "Principles and Lessons in Packet Communications." Proceedings of the IEEE, Vol. 66, No. 11 (November, 1978), pp. 1320-1329.
- LAP76 Lapidus, Gerald. "Line Monitors Open a Window on Data Channels." Basics of Data Communications. Edited by Harry Karp. New York: McGraw-Hill, Inc., 1976.
- LES78 Lester, Vic; Thacker, John, Jr.; and Partington, Donald. "Tech Control for More Uptime and Less Friction." Data Communications, May, 1978, p. 59.
- MAR76 Martin, James. Telecommunications and the Computer. Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1976.
- NAC77 Software Tools for Network Design and Analysis. Great Neck, N.Y.: Network Analysis Corp., 1977.
- OCC77 Occhiogrosso, B., et al. Performance Analysis of Integrated Switching Communications Systems. National Telecommunications Conference, 1977.
- PAR76 Parker, Donald W. "Network Management Centers Help Keep the System Going." Basics of Data Communications. Edited by Harry Karp. New York: McGraw-Hill Inc., 1976.
- POH75 Poh, S.S. Comparison of PWIN and Arpanet Performance Measurement Capabilities. McLean, VA: Mitre Corp., February, 1975.
- PYK77 Pyke, Thomas N., Jr. "Assuring User Service Quality in a Distributed Computer Network." Tutorial: Distributed Processing. Edited by Dr. John Berkowitz. New York: IEEE, 1977.
- RUD78 Rudowsky, I. and Rootenberg, J. "An Upper Bound for the Total Service Time in Packet Switched Store-and-Forward Communications Networks." Computer Networking Symposium. New York: IEEE, 1978.
- SMI78 Smith, Jonathon P. and Moulton, Pete. "Telenet, Tymnet: What the Difference is." Data Communications, October, 1978, p. 67.
- STI79 Stiefel, Malcolm I. "Network Diagnostic Tools." Mini-Micro Systems, March, 1979, p. 63.
- TES80 MS-109 User Documentation, Volume 2, McLean, VA : Tesdata Systems Corporation, March, 1980.
- TOB78 Tobagi, Fouad A., et al. "Modeling and Measurement Techniques in Packet Communication Networks." Proceedings of the IEEE, Vol. 66, No. 11 (November, 1978), pp. 1423-1445.
- WEC76 Wecker, Stuart. "Packet Switching with Assorted Computers in a Private Network." Basics of Data Communications. Edited by Harry Karp. New York: McGraw-Hill, Inc., 1976.
- WIN74 PWIN Users Handbook. Reston, VA: Joint Technical Support Activity (now the Command and Control Technical Center), July 1974.